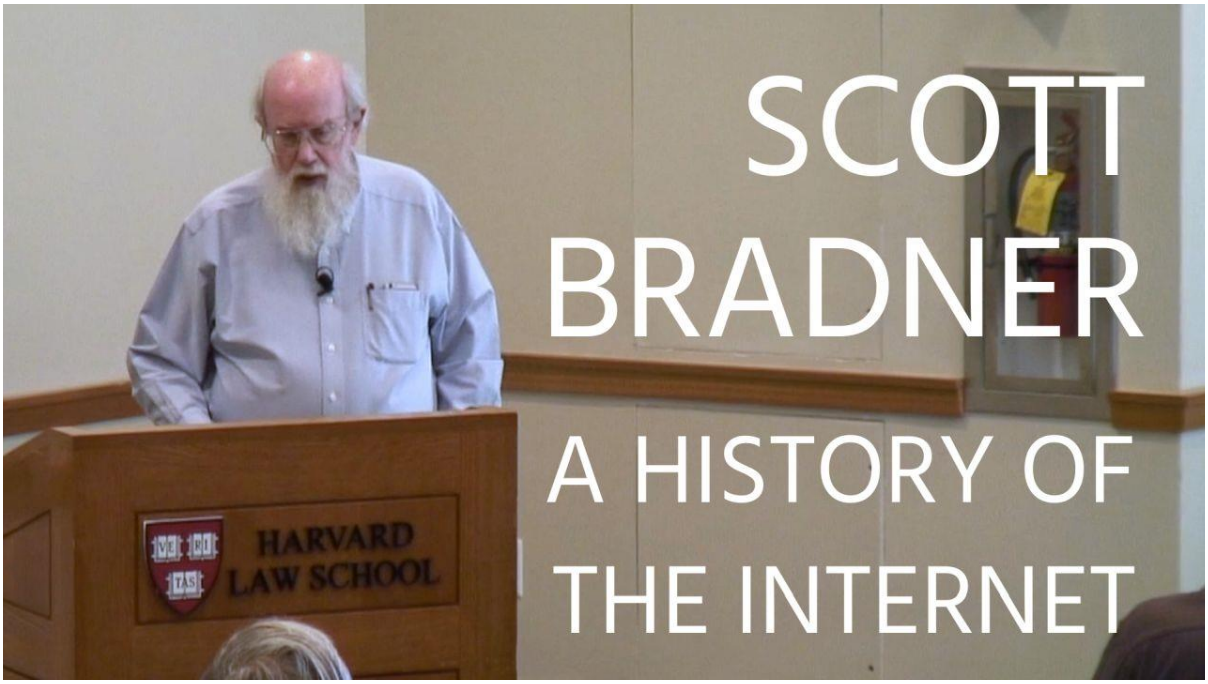

Scott Bradner - A History of the Internet

Berkman Klein Center, February 5 2019

Scott Bradner: Okay, so I'm going to talk about a history of the Internet. This is not the history of the Internet. There isn't one. This is a very much Rashomon kind of environment. I've talked to a lot of people that were involved in this. I lived through part of it.The people that I talked to, some of their stories have changed over time, so even within one person, they have different memories. But, I'm going to talk about a sort of an overview, episodic overview, this is certainly not a complete picture, but it's a basic picture.

It all started with a grapefruit, the electronic grapefruit, Sputnik. When the Russians popped up Sputnik, the U.S. military industrial complex and the political complex went apeshit. Lyndon Johnson said "The skies are alien". Well, Lyndon had a way with words. But it was a big deal.

Luckily, at the time, we had a president that actually believed in science. Imagine that. A person that believes in science. Well, Dwight Eisenhower, being a military guy, definitely believed in science. And he, along with his science advisor, put together in three months an Advanced Research Projects Agency for the Department of Defense.

The idea of putting together anything in Washington in three months is amazing, but he did it. He got it going. The Advanced Research Projects Agency. Basically, exploring new stuff was the formal mission. The informal mission -- anything in italics here is quotes from source material -- the actual mission, according to contemporary materials, was to prevent technological surprise like the event of the launch of Sputnik.

So, here's the ARPA setting up process, starting to work on things. But, it wasn't just ARPA that was being funded by the U.S. government to do work. RAND was funded. The Air Force was funding RAND and Hughes Aircraft, and there's a fellow named Paul Baran that was working there. Paul Baran came up with the basic concept of what is today's Internet. In a 1960 paper, he described ways that you can build a reliable network out of unreliable components, so cheap inexpensive devices we call routers today. You put enough of those things together with enough redundancy, that can be a reliable environment.

And then, in 1962, his seminal paper, followed up two years later by an 11 volume set, described the full set of Internet protocols that we have today. The full set of conceptual Internet protocols. Except, it wasn't Internet, it was a network protocol.

So, he came up with the idea of breaking communications up into chunks. He called them standardized message blocks, but the idea was you had a big file, you made it into little pieces. You had a voice conversation, you made it into little pieces. Each of those little pieces had an address at the beginning, said where the piece was going, an address after that saying where the piece is coming from, some other information, and a payload, which was a chunk of voice or the chunk of data.

It was a distributed network. The bottom picture is the distributed network. The little dots are the routers, the forwarding engines. All the messages, the blocks, were store and forward, so they'd be forwarded to the next little router along the path, saved there, figured out where to go next, and then go next. Had the advantage that if you took out some of those dots, there were other paths around it, and the routers would figure out those other paths and forward the traffic along the other paths.

This is indicative of the times. This is a chart out of his 1962 paper, which is plotting the probability of a message getting through against the probability of nodes being destroyed in an atomic attack. So, you take out half the nodes. If you got enough nodes, the message will still get through. So, he was describing a resilient infrastructure to survive an atomic attack. Big advantages to the message block based thing, the packet based thing. You can put together multiple communications on the same wire. They can be running at different speeds. If there's a failure, if somebody blows up a router, or a router dies, or somebody hits a phone pole and takes out a fiber cable. It takes a minimum time to switch over to a new path. Many different applications, including voice. He described using it for voice in 1962.

The environment he was working in was the Cold War. It was a pretty ugly environment. He was very much anti-war. In his oral history, which is pointed to here, he talks about what he was trying to do was create a command and control network which would survive a first strike.

But, unlike most people, he wanted the Russians to have it too. He wanted it so that the Russians wouldn't think that the U. S. wouldn't retaliate from a first strike, and he wanted the U. S. to not think the Russians could retaliate. So, he wanted both to have it, it was primary. All of his documents were open, they were public. There was two that weren't, to do with security of the system. But, basically, his idea was to make a public environment which keeps both sides from killing the other.

It didn't work. I mean, it didn't get built. He describes in his oral history of going into AT&T to tell them about this, and after he made his presentation, they said, Son, let me tell you about how telephones work, and spent the next few hours telling him how a telephone works, and at the end said, do you understand now, and do you understand why this won't work? And he said, no.

He moved on. He was offered an opportunity to get it built, but the Defense Communications Agency would've been the ones that had to build it, and there was no conceivable way that the Defense Communications Agency could have built it. He realized that they didn't believe in it, and they would've built something which would've failed, and then it would kill the whole concept for eons. So, he declined to have them build it.

That was him.

Meanwhile, J. C. R. Licklider at MIT was talking about intergalactic networks, which was really talking about global networks, and talking about how people would interact with computers in the future. He wasn't the first to do that, but he was an important one to do that. In 1962, he argued for, created, and then started up the Information Processing Techniques Office in ARPA. By the way, this was called ARPA, Advanced Research Project Agency, some of the time, and DARPA, Defense Advanced Research Project Agency, some of the time. Congress flipped its name from time to time for congressional reasons, nothing to do with logic. But, he created this IPTO, and IPTO will come along because that's the organization that built the ARPANET. Bob Taylor took over IPTO in 1965, and got the money, got a million bucks, to build what became the ARPANET.

Now, it didn't cost a million bucks, it cost a few million bucks to build the original ARPANET, but, still, he got the initial money to do it. He was a pretty important person in this, because he believed in the concept, he believed in what Licklider had been talking about, and he believed that networks and interconnected computers were an important thing.

Meanwhile, over in England, Donald Davies was working on the same general concepts that Baran had. He developed them independently. In a paper that he wrote later on, he says Baran invented the whole thing, he invented more details than Davies had. Davies came up with the word packet instead of standardized message block, which doesn't roll off the tongue anywhere near as wel, so, packet is a nice thing. He was introduced to Baran's work, and then he said okay, that's it, and learned from Baran, but also taught Baran.

Next person in this picture is Larry Roberts. He was appointed by Taylor to actually build the ARPANET. A key thing, the idea here of the ARPANET was to be sharing large computing resources. The Defense Department and other government agencies were buying big computers for people. Big computers for researchers, not for people, for research organizations -- and the computers in those days were very big, very expensive, and there weren't many of them, so you want to share them, ARPA wanted to share them. Licklider and company said this is the right way, we share these resources.

The people who owned those resources, they weren't all that interested in sharing them, but they were persuaded by the fact that the guys providing the money said, you will share. But, they really didn't want their computers to be doing the processing of sending data around. They were okay if somebody wanted to remotely log in and use it, but they didn't like the idea that their computer was going to be processing all those packets that went through.

Wes Clark came up with the idea of building separate little boxes which would do the processing, and that rescued the basic concept, because Roberts wasn't going to get through the idea that they were going to have sharing resources, where 20 percent of the machine was going to be taking up with moving packets around. So, Clark's idea of having a separate box, which is now called a router, was a key thing.

Roberts had come up with the idea of building such a network, but he didn't know how to do it. He didn't know what technique, what things to use to do it, what technology. He knew he didn't want circuits, because circuits were going to be dedicated between links, and he wanted to share those links among multiple researchers, so he didn't really know what to do.

He went to an ACM meeting in Gatlinburg, Tennessee, 1967. He presented a paper about the justification and the general concept of this ARPANET, this resource sharing network. Another person at that same meeting was a fellow from Don Davies group who presented about packet networks.

Roberts told me, personally, that he had never heard of packet networks before this meeting. He had never heard of Baran's work and he had never heard of packet networks. Scantlebury told Roberts, in the meeting, about Baran's work.

Roberts had volumes of Baran's work on his shelf in his office, and he went back and read them, and he said that was a revelation. He basically took all of Baran's ideas, modified by Davies, updated by Davies, and that's what he decided to use, he was going to use packet networks for the ARPANET. And, he took over the IPTO in '69

So, here's the ARPANET, starting up. In '68 an RFQ went out, 140 companies were asked to bid on building these little message processors, what turned into routers, and only 12 companies actually bid. IBM did not bid. Later on, they justified it by saying they didn't think it could be done cost effectively. AT&T didn't bid, but they said they didn't bid because it wouldn't work, packet networking was never going to be useful for anything, just like they had told Baran.

BB&N out here in Fresh Pond won the contract, and within less than a year, they delivered the first router, the first IMP, Interface Message Processor, even though Ted Kennedy called it the Interfaith Message Processor. Entirely different concept.

So, the first node was delivered to UCLA in '69, and a month and a half later, the first message was sent between two different IMPs, in two different locations. Initial networks, four nodes, Kleinrock and others were there. It quickly expanded. By next year it was on the East Coast. John Klensin reminds me that MIT got on before Harvard. So, here it is. There you are, John.

And then, BBN got in there someplace, and, by '73, it was international. By '80, it was a couple of hundred hosts and 20,000 users. But, those users were people in the U. S., they were only people with U.S. government funding, or who were staff on computers that were funded by the U. S. government, it wasn't open to all people. I mean, I got my first account in '71 by conning my way into being staff on computers. The PDP-10, that was the ARPANET interface in the Aiken computer labs, otherwise I wouldn't have been able to get on because I wasn't government funded. I was, time to time, but not at that particular time.

Everything changed in '83 when CSNET made a deal with the ARPANET that any organization getting their CSNET connectivity, Computer Science Network connectivity, via the ARPANET could open up that pipe for every faculty, staff, and student, for email -- not for file transfer, but for email. So, starting in '83, we started creating, teaching generations of students about this networking stuff. Before, it was only the computer scientists, but now it's generations of students, so that, when they graduated and went out to business -- We need this, join up!

So, that was a huge forcing factor of getting the Internet dispersed.

So, basic question. Baran's work was for anti-war, it was for protecting against the bomb. Robert's work was for sharing large resources. I've talked to a number of the people that were the worker bees at the time, and they all say, from their point of view, it was sharing resources, it had nothing to do with the bomb, it had nothing to do with resilient infrastructure for the bomb. Yes, the resilient infrastructure was there. It's like what Bruce Davies said, he was building his network over the French telecommunications system, so he knew it was unreliable. It was a feature.

And so, he had to build a network that would not depend on reliability underneath. The packets could be lost, they could be duplicated, they could be reordered, and that was okay, because the network was controlled from the edges. The source and destination machines took care of making it reliable, took care of reordering things, took care of asking for retransmission if things were lost. So, the worker bees at the time said, we're building a resilient resource sharing network.

Steve Crocker said that this was what he believed. He was one of the people in there. But, he, later on, a few years ago, talked to one of the higher ups in the Pentagon at the time, and the guy said, you don't know what I was telling my boss.The money was coming from the Pentagon. The money was coming from the Defense Department. So, it's very conceivable that the top levels of the Pentagon were telling Congress, and they were telling each other, this is for anti-nuclear, this is for surviving nuclear war, even though the people that were doing the work didn't think that.

So, the answer is ambiguous, and it's not something I can answer. I can only present you with the information.

So, now, Bob Kahn shows up in this picture. He got a Ph.D. at Princeton, but then went off to work for BB&N and worked on the IMP, and moved to ARPA IPTO, and then made a big demonstration, in 1972, of what was then the ARPANET sharing resources, in Washington.

One of the people that were demonstrating there, I think it was Bob Metcalf, said that he was demonstrating all day of how to do resource sharing, and it was working perfectly, and a bunch of people from AT&T came in late in the afternoon, all dressed up in their pretty little suits, stood at the back of the room, and while he was demonstrating, it crashed for the first time, and he turned around just in time to see them all grinning and laughing, and he said, it was at that moment that I realized they were the enemy. They didn't really believe in this stuff.

So, Kahn is there, we'll come back to Kahn.

Danny Cohen was a Harvard student and then MIT student. He was doing flight simulators work between Harvard and MIT, and using the ARPANET link for that. He was mostly working, later on, on speech, and he did do speech work over the ARPANET, but he realized that the original ARPANET protocol, NCP, Network Control Program, was a reliable protocol, and a reliable protocol is reliable because it knows when something's missing by timing out, and saying, oh, I need that packet again. Well, if you're in the middle of speech and you get a two second dropout of the middle, it doesn't work very well. So, he said, we need something which is not part of the reliability picture, and he was given permission to use an unreliable pathway through NCP in order to get his speech work done.

Over in France, Louis Poussin was working on his version of a packet network, peer datagram network, reliable over France Télécom lines, so it couldn't possibly be reliable, all of the stuff end-to-end.

Bob Kahn asked Vint Cerf to figure out a way to be able to make a scalable network.

The ARPAnet as it began was to connect computers together, connect a computer at Harvard to a computer at Berkeley to a computer at MIT. It was not to connect networks together. Kahn realized that scaling wise you needed to connect networks together, and he worked with Vint Cerf to come up with a set of protocols to interconnect networks, and the computers on those networks, and that's the Internet Protocol. The Internet Protocol, layered protocol.

The first version of this was a reliable protocol called ITCP. It was meant to be reliable because, of course, you want a reliable transfer of data. Danny Cohen and John Postel and David Reed cornered Vint Cerf in a hallway in Southern California and said, we can't do this, we need a way to do something that's unreliable for voice and for management. If your network is all screwed up, you don't want to have forced reliability on your management protocol, which is talking to your routers, because the message will never get through, so you want an unreliable method. So, they added what is called UDP, User Datagram Protocol, otherwise known as Unreliable, that enables the Voice over IP that we use today.

The general idea that the network is just going to deliver stuff, and is not going to be involved in the value decisions, was articulated in the end-to-end arguments and system design from Saltz, Reed, and Clark from MIT. Basically, you don't put stuff in the network which is redundant with what the end-to-end systems can do. The end systems knows what service they want out of the network, the network cannot know that. There's no way for the network to figure out what service they need, what level of reliability, what latency, all of that kind of stuff, they can't figure that out.

But, the telephone companies and the carriers thought it was a really good idea to have the network understand what's going on. That conflict, between a stupid network that just delivers the bits, and an intelligent network -- the telephone company called their network the 'intelligent network', because all of the smarts were in the telephone switches. Star 69 was a telephone switch thing that took 10 years to deploy.

So, there's an article called NetHeads vs. Bellheads in Wired Magazine. If you've not read that, read it, it's a very good explanation of that conflict. David Eisenberg wrote 'The Rise of the Stupid Network' while he was an AT&T employee. He wasn't one for long, because this message was basically, AT&T, the network, should be ignorant. The network should simply deliver the bits, just get out of the way and deliver the bits. Of course AT&T didn't think that was a good idea, and to this day they won't let David post a copy of his paper on his website, because he was an AT&T employee when he did that. If you go to his website, there are pointers to other people who are posting it.

So, the Internet as it evolved, from the ARPANET on, really didn't need control, didn't need centralization, it needed coordination. Everybody had to agree that this field in this packet, when it said 25, meant email. They didn't care what the number was as long as everybody agreed the same number. So, somebody wrote that down as the 25 means email, that was John Postel. Also have to say who has what IP addresses, and John wrote that down.

When Harvard needed its first batch of IP addresses, I sent John Postel a message saying, we need IP addresses, and he sent me back a message saying, here they are, 128.103.0.0 Class B, that's Harvard's. Harvard's gotten more since then. And then, there's got to be some pointers to, when you resolve DNS names harvard.edu or www.harvard.edu. There has got to be somebody up there pointing at the top level. Where are the computers that resolve the names for .edu? There's a coordination, not governance. They're not control.

All of that stuff was done by John. He evolved it into the IANA, Internet Authority for Names and Numbers. That turned into ICANN a little later on.

So, what did the powers that be think this Internet thing was? It was irrelevant, it's a toy. This is best effort, packet based, no guarantees, no reliability, no security, nothing. It's a toy. IBM, AT&T, and particularly important, the regulators -- this is a toy, it's a research toy, it's a research network and a research toy, so it's not important enough to regulate. That was key, because if it was regulated, you would have to get permission to do things.

The bumblebee is on there because IBM -- I was participating in a lot of IBM user groups at one point -- and IBM said, quote, you cannot build a corporate network out of TCP/IP, direct quote from a fellow named Gray at IBM. I was the co-chair of the TCP project within SHARE, within this IBM user group, we chose a bumblebee for our badge symbol. You put badge symbols like, if you were in communications, an octopus because it goes everywhere, if you were numerically intensive computing, it's a circle with a 42 in it. If you didn't know what that meant, you didn't get one.

We chose a bumblebee because IBM said TCP/IP wouldn't fly. Aerodynamic theoreticians have said that bumblebees couldn't fly. Bumblebees didn't give a crap, they kept flying. We didn't give a crap, we kept going. IP kept going.

So, connectivity evolution. Let's start with the ARPANET, then CSNET, and NSFNET, the regional networks, NEARnet, which was mentioned in the introduction, BARRNET, people like that, and then migrated to commercial ISPs. The Feds were completely out of the business of providing network connectivity for the public in the mid-1990s, and then we got carriers. It moved from something which was Feds, to regional networks, to commercial carriers, to commercial Internet Service Providers whose business was providing Internet service, to carriers whose business was content, and is now doing Internet provision, and that comes back to haunt us in a bit. Then the web expanded, the web exploded, and made it so that anybody could do anything.

I was asked a few years back, looking forward from the 1980s to the Internet of the day, what was my biggest surprise? I thought that the Internet was going to be an important thing in the mid-90s, it was going to be an important thing for everybody, the telecommunications infrastructure. My answer was, Mom surfing. It never occurred to me my mother would use this, because it took magic incantations to do it. The web changed all of that, the web made it so that anybody could use it, and the usage exploded.

So, what made the Internet important? What were the key things? Well, it was the end-to-end model and neutral networks. Network neutrality, the packet, the network just delivers the bits, doesn't get involved in deciding what bits should be delivered which way, doesn't decide whose bits should be prioritized or whose bits should be let through. The networks are neutral.

That allowed permissionless innovation. Anybody, like Tim Berners-Lee, could invent something like the web, or somebody could invent Skype, and you didn't need permission from the carriers to do it. You just did it. You didn't need permission from the regulators to do it, you just did it.

Of course, capacity increase with Moore's Law and the fiber bandwidth was a big deal. We didn't change the net for any particular application. The net was not designed to support some particular thing.

And the IETF, when we first started getting the big push from the telephone companies, they wanted to change the protocol to make it better for voice, and we said, no, the Internet is good enough for stuff. The stuff it's good enough for, use it for. If it's not good enough for it, don't use it. It's a pretty simple equation. Now it's gotten better for everything because of the speed increases.

And, no regulations meant we had permissionless deployment, so, we could innovate and we could deploy.

And then, very importantly, the Communications Decency Act was overturned by the Supreme Court as unconstitutional. That was an act which required you personally to ensure that, if you sent something on the Internet that somebody, anywhere in the world under 18, shouldn't see, they didn't see it. Technically, that's a bit of a tough hoe. The fact that it was technically impossible was irrelevant to the court case.

The Supreme Court decided on it, not because it was technically impossible, but because it wasn't the best way to achieve these results. They thought the best way was to put filters on the individual kids' computers, rather than trying to block the entire world. Because, of course, most of the world doesn't pay attention to the U. S., so, under the First Amendment, this was just too much of a burden on speakers to be able to be supported.

Section 230, on the other hand, of the Communications Decency Act, is the section that permits Harvard to be an ISP, and not be killed if when some Harvard student downloads music. We can be a neutral thing, we're not responsible for the actions of our users, as long as we do due diligence when we're informed of those actions.

And, of course, the key thing that made the Internet successful is that everything is bits.

Everything is bits. Ada Lovelace figured out that in 1843, that computers, analytical machines were what she was dealing with, weren't limited to dealing with numbers, they could deal with anything that those numbers could represent, such as music, and she talked about using it for music.

So, what's the future of the Internet, and technology wise? All of these things, ATM, MPLS, per-flow queuing, and things like that, are things that different people have proposed, different organizations have proposed, to replace the Internet that's here, almost all of those are circuit-based. So, get rid of the unpredictability, put in circuits, that's ATM, MPLS, per-flow queuing, etc.

By the way, how many people have seen the movie Touch of Evil? It's a reasonable number. It's a great film noir film from Orson Welles. This is a scene in there where Orson, as a corrupt sheriff, comes in to Marlene Dietrich as a brothel runner, madame, and says, Tell me my future, Come on, read my future. You haven't any. What do you mean? Your future's all used up.

Every one of those things, ATM, MPLS, per-flow, etc., is some company, some organization, some standards body, some government, has said the Internet's future's used up, we have got to do it right. And the Internet's kept going.

The 3G, 4G, 5G thing here -- first time I heard this was in Taiwan, by one of the telephone people saying, 3G is around the corner, and when 3G shows up, you won't need any local area networks, you won't need any corporate networks, it'll all be provided by the telephone company. Same thing they said for 4G, they're saying that now for 5G. I didn't believe it then, I don't believe it now. And other things...

But, this Internet thing has been impactful. It's destroyed businesses, it's destroyed societies, it was a key facilitator of the Arab Spring, it's a big hit.

This paragraph at the bottom here is from the first international telecommunications convention treaty of 1865, and it was the Telegraph Convention. It was required that telegraph operators watch their traffic, watch the telegrams, and block any that would be in danger of the security of the state, violate the laws of the country, public order or morals, and to report it. So, the very first regulation we have on international telecommunications was a spy one, spy on your users, and report those that are doing something immoral.

That's something we do not have in the Internet. We have been regulation-free. The CDA tried to do that and was killed, in the US. That doesn't mean that it doesn't happen in other countries, it does.

So, what's the future, on the control side? Carriers, of course, have been trying to control it ever since they figured out that there was something there to control, that it wasn't entirely bumblebees and uselessness, it was something that's really there that could actually add value. Some of the big carriers like Comcast are now getting more money from Internet service than they're getting from video.

Governments, of course, the same, petitioning the ITU to take over governance of the Internet so they can control it, so they can avoid the loss of telephone revenue, for example, things like that. Certainly, the current FCC is blowing away the Title II Network Neutrality Rules, are enabling carrier control. They are making it completely legal for the telephone carrier, ISP, to control anything, like what you can talk to, who you can talk to, what you can use to do the talking.

Newt Gingrich was once quoted as saying, a role of government is to guide technology. Not something I'd necessarily want the Congress to do. A guided or controlled Internet wouldn't have been what the Internet has been to date.

22 years ago, or so, I testified in the Communications Decency Act hearing, I was the witness for the American Library Association, and I said that the power of the Internet was chaotic. The power of the Internet was the forum to innovate, which was chaotic. You could not predict what would succeed, and what would fail, it was chaotic, but the fundamental power of the net was represented by that chaos, and the question on the table today: Will that forum continue?

That's it.

So, questions, but use the microphone because the people out there in audio land need to hear. John?

John Klensin: I'm going to skip quibbling about history, and ask a more interesting question. You, in my mind, correctly denounced the efforts to fix the Internet to make it better for particular applications, and similarly denounced some of the related issues about pulling things together and switching toward more circuit like things, again, to facilitate particular applications. How do you feel about changing the Internet to make it better for the web? Or, to put it differently, would you like to comment on QUIC?

Scott Bradner: I think that anything which makes it better for one application, makes it worse for others, and there's no perfect unity. General purpose computers are powerful because they're general purpose, not because they're good at math.

Any other questions?

Participant 1: Can you talk more about Title II and what that has actually affected?

Scott Bradner: Well, the history there is that the Federal Communications Commission, under Powell, came up with four principles of a proper Internet. They were that you could use any device you wanted to, any legal device, you could access any legal content, and you could access any legal application, and you had transparency of what was going on. The FCC put out those four principles, tried to enforce them, they got sued, and they got kicked out, because they had done no proper procedures to adopt those policies.

They then went through a procedure to try and adopt something very similar. They got sued again, and it got thrown out because they didn't undergo the right process. They didn't have the statutory authority to do it. They tried again, and it got kicked out again, and that last time, the court said, the only thing you can do, in order to be able to do this kind of regulation, is to declare that the Internet service is a common carrier, so it's under Title II of the Communications Act.The FCC did that. They got sued and the FCC was upheld in court, that it was legitimate for them to decide that Internet service providers were common carriers, and therefore had to abide by the rules of common carriage.

The particular thing the FCC did was say they're subject to Title II, but we will only actually enforce, will forebear on other, rules, that are relative to those four principles. So, the 200 other rules out there, for who you can give free tickets to the movies to, literally, that's the kind of thing that's in the telephone regs, they don't apply.

A lot of people really didn't like the fact of Title II, because it gives the FCC so much power. You have to pre-get how much we're going to pay for things, the tariffs have to be pre-approved, the application has to be pre-approved. All of that kind of stuff is potential, but the way the FCC approved this particular bill is they're going to forbear on all of that.

People don't like that because it's got the potential there. Somebody could sue the FCC later on and say, These people are being unfair, under Title II, you can do this, go do this. And, the courts could, in theory, tell them to do something, like control prices. but, the courts had said the only way they could do it is Title II.

So, the new FCC overturned that, threw out that Title II stuff. But, in their new regulations, say the ISPs can do anything they want, except they have to be transparent about it. The rubric under which they said they have to be transparent is exactly the same rubric that the courts threw out, when the FCC was trying to do the four principles.

So, it's in court today, it was argued on last Friday, it could be overturned on a number of different reasons. It could overturn the overturning of Title II, because the process, the FCC used was bad. They literally ignored 20 million comments, they completely ignored them. Now, it happens that a significant chunk of those comments were done by bots. Research later on showed that all of the bots hated network neutrality, and wanted Title II overturned, and all of the humans wanted it retained.

Gives you some information about bots. But, so it could be overturned on that basis, it also could be overturned on the basis that they're imposing this transparency requirement that they don't have any authority to impose. It's potentially very big blunderbuss, but it's the only one the courts would let the FCC use.

Participant 2: The packet design of TCP/IP is necessarily somewhat arbitrary, I mean, you could design another packet that would still work, and you used to have DECnet and BITnet and, you know, Joe Blow's net. So, is that packet design somehow superior, or is it entirely arbitrary?

And, if we can design a better one somehow, not whether, when will it be installed to completely rewrite the Internet?

Scott Bradner: Well, first of all, TCP/IP is as much a regulatory accident as anything.

The Defense Department agreed to switch the ARPANET to TCP/IP on January 1st, 1983. The NSFnet, when it started up, Dennis Jennings said, only IP is enabled, it can be run over the NSFnet. You couldn't run DECNET. Harvard had a DECNET computer, talking to a DECNET computer at Princeton, over an NSF sponsored regional network link, and the physicists at Harvard said, We would like to use DECNET. Is that okay? And NSF said, No, you have to use TCP/IP. So, it enforced that.

Scott Bradner: Second thing is that ARPA paid Berkeley to take the BB &N TCP/IP code and put it into Unix, and then startups could buy a tape with all of Unix, including a full TCP/IP stack, for 50 grand or so, much cheaper than it would be for them to build their own.

So, the Department of Defense basically pushed out IP to everybody much cheaper than building your own, so that meant that all the little startups had it. The big guys, they started moving to OSI, like DECnet tried to move to OSI. It was so much more complicated, and so much more fragile, that it, just didn't work.

Finally, the U.S. Government, which had said, if you're going to sell networking stuff to the U.S. Government, you have to support OSI, backed off and said, you can buy whatever would do your job, including specifically TCP/IP. So, it's somewhat of a regulatory accident that's it's IP.

DEC was out there with DECnet, IBM was out there with SNA, Xerox was out there with XNS, Novell was out there with IPX, they all had their strengths and weaknesses. IP is there because it became the common language.

Participant 3: Do you think it'll ever be replaced by something which is somehow packet wise superior?

Scott Bradner: So, the ARPANET switched from NCP to TCP/IP officially on January 1st, 1983. It took 6 months. There were a couple of hundred computers. If you project that out, if you wanted to switch the Internet of today, which has a few billion computers, to a new protocol, it would take a few hundred thousand times the projected lifetime of the universe to do, so don't bet on it.

Participant 4: So 5G. I ran into a, a telecom guy at a party the other night who said that the Internet was a Trojan horse for the telecoms, and they've been trying to get the horse out ever since. 5G is how they think they're going to do it, and he's a telecom guy. The idea, he said, is you're going to get very high speed, you're going to get it on your phone, you get it on everything, it's going to have local storage, and they're going to compete with Amazon by blowing up their cloud and distributing it everywhere, and they're going to have their cloud, and they're going to charge. They're going to get in on Hollywood's business, they're going to bid on Netflix business, and it's going to be all top down distribution, and then they finally win.

So, two questions. One is, is that true, or close enough, and B, how do we fight that shit?

Scott Bradner: Well, first of all, if it was a Trojan horse, it was against the telecoms.

Participant 4: Yeah.

Scott Bradner: So it wasn't for the telecoms.

Participant 4: Oh right, that's definitely true.

Scott Bradner: Yeah, and that's the same pipe dream that the carriers have had from the beginning. There was a whole big todo in the late 90s that content was king, that, somehow, when the carrier started providing you with the ability to haveTiVo on demand, a couch potato heaven of TiVo, you're going to triple or quadruple your monthly fee to the telephone company.

Well, content isn't king, and, in the particular days when the telephone companies were saying that, there was ten times more revenue from telephone calls than there was from all of Hollywood put together, and yet they thought that, once they got the movies, they would increase their revenue by five fold.

It didn't make any sense then, it doesn't make any sense now. You really want to depend on Verizon for your entire life experience? I don't think so. It's not quite as bad as what a fellow from MIT said a few years ago, when we used to have NYNEX used to be the telephone carrier in the Boston area, and he said that that was the Iroquois word for moron.

Verizon isn't quite that bad, but it's not good.

Participant 6: Hi, the story from the party sounds sadly like the plot of the TV show Silicon Valley, which is satire, so I'm not sure I would treat it with a great deal of respect, but my real question is, whatever happened to IPv6 and IPsec?

Scott Bradner: What do you have in your hand? That one.

Participant 6: IPhone.

Scott Bradner: You got IPv6. IPv6 is native on the iPhone. It's native on most smartphones. It doesn't have to do anything right here, because Harvard doesn't support IPv6, but, if you're using the Verizon network…

What network are you using?

Participant 6: T Mobile.

Scott Bradner: I'm not sure about T-Mobile, but Verizon supports IPv6. Comcast does IPv6 to your house. The big thing with v6 was, it was invented to solve a particular problem, which was running out of IP addresses, and try to do a little bit more at the same time, but it didn't provide a huge number of carrots, so it didn't have a lot of reason to move there unless you're running out of addresses. We're running out of addresses, so big carriers like Verizon and the like, Comcast, can't use v4 addresses because there ain't enough of them, so they're using v6. It's transparent to you because it's infrastructure.

Just like Mom surfing, it's completely transparent to you. But it's there, and it will continue to expand because there isn't any choice. There's no v4 addresses available, not in the U. S. anyway.

Participant 6: And IPSec?

Scott Bradner: IPSec? That's VPN. Go to your system controls on the iPhone and press VPN,

that's IPSec.

Participant 7: Hi, so you at first mentioned that regulation would be a bad thing, but then you seem to agree that things like net neutrality were positive, so would you agree that there might be some types of regulation that could be positive, and what would those types be? What are we lacking right now in terms of regulation that perhaps promotes or enforces openness?

Scott Bradner: Well, at this instant, what we're lacking is the enforcing factor to be sure that the ISP keeps out of the way, that's what Nat Neutrality is all about. We had that with the Title II stuff, we don't have that anymore, because that was overturned, though it's in court, so we'll see what happens there.

That regulation is not telling you to do something, it's telling you not to do something, and that's sort of different. The regulation that telephone companies undergo, for example, is, I want to bring up star 69. Well, I've got to get permission, I've got to go through a regulatory process to say how much that's going to cost, and then that's never been the case for the net.

But, with the ISPs now in the middle, it's not a regulatory thing. The ISP can say, well, if you subscribe to this service, it doesn't count against your byte limit, your download limit, but if you subscribe to this other service, which we don't have a business relationship with, then it does, and now the ISP is in the middle deciding who's winners and who's losers. So, a regulation that says, Keep out of the way, is very different than a regulation that says, Ask me first.

At one point a while back, Verizon wanted to regulate so that all Voice over IP had to go through central servers. Now, they said it was so the FBI could enforce wiretapping, but really it was so that they could bill for it. That kind of regulation to force that would inhibit the innovation, the ability to have that chaotic environment to innovate. Anything which gets in the way of that, asking permission first to innovate, or permission to deploy, is the inhibition of the dynamics which got us the Internet.

John, again?

John Klensin: I want to give a slightly different answer to that last question, because I think another aspect of that is important. One of the things we have been lousy at, as a society and a networking community, is distinguishing between things which are ultimately legal matters and things which are ultimately technology matters.

We try to substitute the technology for the legal issues. We try to substitute, much more rarely lately, but more so in years past, legal solutions for technology issues. We're spending tremendous resources on the network, usually at the wrong end, trying to deal with spam and phishing and other kinds of attacks, which are basically legal problems.

And, if you think about those as part of the regulatory picture, the answer to the answer Scott gave you to your question is we better get a lot better at figuring out the difference between the legal problems, the criminal code problems, and the technical solutions, real or alleged.

The question I wanted to ask you is that one of your slides, you put up a picture of the logos of a large collection of relatively early entry commercial ISPs...

Scott Bradner: And then a few more recent ones.

John Klensin: And a few more recent ones. But, in the last few years, we've seen significant consolidation there, to the point that, if that picture were drawn today, it would have a lot fewer icons on it.

Scott Bradner: It basically has the bottom three, I think.

John Klensin: Yeah, do you think that's an issue going forward? And if not, why not? And if so, what do you think we do about it, if anything?

Scott Bradner: Well, I think it's a significant issue, not in and of itself. It's the pipe dream, the fiction that the FCC undergoes, saying there's competition for the residential market. There isn't.

And yet, the FCC says that 60 percent of the country, or 70 percent of the country, has competition, realistic competition for high speed Internet service. Their definition of high speed is 5 megabits, so that's questionable there. 5 megabits down and, I think, 1 up, which is piddling in Europe and Asia.

I think that fiction is a very serious issue, which is the fiction which is driving this getting rid of Title II, it's the fiction that's driving the Congress in general. A few years ago I wrote a column that, the people in Congress think that there's competition on the

Internet provider environment, but they also think that the driver of the Internet is the carrier.

This is like you've been asked to evaluate highways. You're standing beside a highway. But, your eye is in your ankle, and all you can see is the asphalt. You can't see the cars and the trucks and everything going over the highway. So, your answer to make more highways, is to enable asphalt manufacturers, not to enable car manufacturers. The Congress, its eyes are in their ankles, all they're seeing is the telecommunication carriers. They're not seeing the fact that those carriers are out of the way, drove. That's the Googles and the Amazons and the Skypes and everything else that's riding over on top. That, they can't see. They don't see that that's part of the Internet, and that's, to me, the biggest problem.

No competition, the myth that there is, and belief that the Internet equals carriers. It's just crazy.

Well, on that positive note...